

Up and Away: A Cheap UAV Cyber-Physical Testbed (Work in Progress)

Ahmed Saeed[†], Azin Neishaboori^{†‡}, Amr Mohamed[†], Khaled A. Harras[‡]

[†]Department of Computer Science and Engineering, College of Engineering, Qatar University

[‡]School of Computer Science, Carnegie Mellon University Qatar

ahmed.saeed@qu.edu.qa, azin.neishaboori@gmail.com, amrm@qu.edu.qa, kharras@cs.cmu.edu

ABSTRACT

Cyber-Physical Systems (CPS) have the promise of presenting the next evolution in computing with potential applications that include aerospace, transportation, robotics, and various automation systems. These applications motivate advances in the different sub-fields of CPS (e.g. mobile computing and communication, control, and vision). However, deploying and testing complete CPSs is known to be a complex and expensive task. In this paper, we present the design, implementation, and evaluation of *Up and Away (UnA)*: a testbed for Cyber-Physical Systems that use UAVs as their physical component. *UnA* aims at abstracting the control of physical components of the system to reduce the complexity of UAV oriented Cyber-Physical Systems experiments. In addition, *UnA* provides an API to allow for converting CPS simulations into physical experiments using a few simple steps. We present a case study bringing a mobile-camera-based surveillance system simulation to life using *UnA*.

Keywords

UAV Testbed, Cyber Physical Experiments

1. INTRODUCTION

Cyber-Physical Systems (CPS) have the promise of presenting the next evolution in computing by bridging the gap between the virtual world and the physical world. CPS potential applications include aerospace, transportation, and factory automation systems [9]. Advancements in cyber-physical systems research is supported by development in other research fields including mobile computing, embedded systems, computer vision, control, and communication. While this development ultimately contributes to advancing different components of cyber-physical systems, most of the work is evaluated either through simulations or through experiments that focus on the specific sub-problem. Furthermore, the complex and multidisciplinary nature of full scale experiments of cyber-physical systems makes the task

even harder especially when it comes to testing individual "cyber" or "physical" components. From a different perspective, current generic CPS testbed [3] are typically complex and incur large deployment cost.

In this paper, we present Up and Away (*UnA*), a testbed that allows for low-cost, rapid development of cyber-physical systems experiments. *UnA* abstracts the "physical" components of the system to allow for experiments that focus on "cyber" components with real physical interactions. We chose UAVs as the physical component of the system due to their deployment flexibility and maneuverability. Furthermore, UAVs are suitable for many application (e.g. surveillance, transportation, and communications) allowing for cyber-physical experimentation with realistic objectives.

Several research groups have recently developed quadcopter testbeds [4, 6, 7, 11]. The GRASP testbed [7] uses off-the-shelf high-end Ascending Technologies Hummingbird quadcopters to demonstrate multirobot control algorithms. The ETH Flying Machine Arena [6] uses modified Hummingbird quadcopters to demonstrate several acrobatic and athletics control maneuvers. A testbed was developed by University of Colorado, Boulder for testing ad-hoc networking scenarios [2]. The testbed uses UAVs which were deployed in a 7 km² area to test different operation conditions of the DSR protocol. The work in [4] shows the development of a low-cost UAV testbed using Goldberg Decathlon ARF model airplane. The work aims at showing the design and development of hardware and software components to produce low-cost UAVs.

Unlike other UAV-based testbeds, *UnA*'s architecture is more concerned with abstracting the control of UAVs rather than improving it. Furthermore, *UnA* presents an architecture that provides default components and facilitates replacing any of them to test its impact on the overall performance of the system. To that end, AR Drones were selected as the UAVs to be used. We preferred using quadcopters to fixed-winged airplanes due to their maneuverability especially in highly constrained spaces (e.g. indoors) including their ability to hover. Moreover, AR Drones are extremely cheap¹ UAVs that come with a large array of sensors and a widely supported open source API to control the drone and obtain sensory information for its sensors (the specifications of the AR Drone are better described in Section 2).

2. UNA ARCHITECTURE

¹An AR Drone 2.0 costs \$300 while an AscTec Hummingbird costs around \$4500 .

In this section, we introduce the design and implementation of *UnA*. We start with our design goals followed by the architecture description. The, We describe the details of how *UnA* handles targets and UAVs localization, UAVs control, CPS processing and communication.

2.1 Design Goals

The development of *UnA* is motivated by the need for a low-cost, easy to deploy experiments for CPS. Thus, *UnA*'s architecture is designed with the following goals in mind:

1. *Multidisciplinary nature of CPS experiments:* The development of CPS experiments requires awareness of several fields including hardware design, control systems, software systems development [9]. This poses significant challenges for developing realistic experiments for such systems. *UnA* was designed with this unique feature of CPS in mind and therefore the *UnA* modules are extensible to support development of any of the CPS's components.
2. *Support of large scale experiments:* CPS experiments with a large number of physical nodes requires the testbed to support different communication models between the nodes. Furthermore, the processing power required to optimize the overall state of the system can grow exponentially with the number of nodes. The *UnA* architecture supports Ad-Hoc communication using AODV for routing to facilitate communication between the different UAVs. Furthermore, it supports off-loading processing tasks required by the UAVs to central or multiple computing nodes.
3. *Seamless control of the UAVs:* *UnA* aims at reducing the complexities in CPS experiments design incurred by the automatically operated mechanical components of the system by abstracting the control algorithms and hardware design which allows for more focus on the "cyber" problems. Furthermore, autonomous control of off-the-shelf UAVs is a difficult task [6, 7]. Thus, *UnA* relies on the easy to use AR Drone API and extends it to use way-point navigation instead of controlling the rotation (i.e. roll, yaw and gaz) of the UAV.

2.2 Overview

Figure 1 shows a summary of the *UnA* architecture. *UnA* allows the UAVs that can communicate with each other, process sensory information locally and act accordingly (i.e. control itself). It also provides a base node that monitors the behavior of all UAVs. Moreover the central node can handle processing intensive tasks that UAVs off-load to it.

The *UAV Node Modules* are responsible for obtaining and processing information about the drone and environment's state, calculating the drone's corresponding objectives accordingly, and controlling the drone's motion to achieve these objectives. *UnA* is built on top of the Parrot AR Drone 2.0 [5]. The AR Drone is a quadrotor helicopter that are electrically powered with two camera: a front 720p camera with 93° lens and a vertical QVGA camera with 64° lens. These drones are controlled through an external computer through WiFi and are equipped with an on board ARM Cortex A8 processor with 1 Gbit RAM that runs a Linux 2.6.32 Busybox. The on-board computing machine is responsible for collecting and reporting the state of the drone to an external computer that controls the drone. The AR Drones are relatively cheap, easy to program and comes with an SDK that provides assisted maneuvers which satisfy most of our

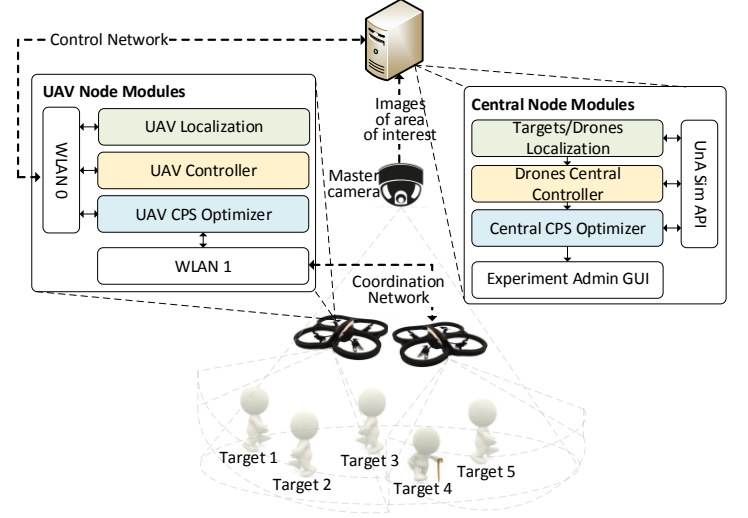


Figure 1: *UnA* Architecture

design goals.

The *Central Node Modules* are responsible for monitoring the state of each drone, manually controlling the drone, and providing part of the environment state information to the drones. Due to the limited processing capacity of current off-the-shelf UAVs, *UnA*'s architecture enables off-loading some of the processing tasks from the UAVs to the central node to enable real time responses to changes in the environment. The *Experiment Admin GUI* displays the state information of the different UAVs. It also supports the manual control of the UAVs for emergency cases. The main purpose of *UnA*'s architecture is to support the rapid development of CPS experiments that use UAVs as their Physical part. While all *UnA* can be extended and customized for different CPS applications, we assume that simulation code was already developed by the experiment administrator. Thus, the central node supports communication between its different modules with other applications through sockets. The *UnA Sim API* provides the appropriate wrappers to allow the simulation code to interface with *UnA* to bring the simulations to life by off-loading all the *UAV Node Modules* functionalities to the simulation engine through the central node.

For the rest of the section, we explain the different layers of the *UnA* architecture according to the color coding in Figure 1.

2.3 UAV and Target Localization

UnA provides two different approaches for obtaining location information of the UAVs: 1) Parrot Flight Recorder GPS module and 2) using a master camera mounted at a high point to track the location of the drone. While the first approach is straight forward, it doesn't work within controlled indoor environments. Thus, computer vision is used for tracking the UAVs in *UnA*.

Multiple UAV tracking using computer vision with multiple cameras is a tricky task to implement [8]. For that purpose, we propose a simpler approach that uses a single camera, mounted at a high point in the center of the area of interest, to track the drones by tagging each drone with a

distinct color tag². Tracking objects of a certain color could be easily implemented using OpenCV [1]. This approach is further extended to track other targets as well.

While earlier work [6, 7] relies on Motion Capturing Systems (MCS) for tracking UAVs. We chose the single camera approach as a much cheaper alternative that provides reasonable accuracy (Section 3)³.

2.4 UAV Control

UAV Controller and *Drones Central Controller* are responsible for moving the drone to a specific X, Y coordinates with a certain orientation within the area of interest. While quadcopters can move in any direction with any orientation, we use a simple controlling method: at the beginning of each command the orientation of the drone is set to 90° w.r.t. the x-y plan of the area of interest. The change in orientation uses the compass for feedback. The UAV then moves the drone along the x-axis to its specific coordinate then along the y-axis and finally it's rotated to the specified orientation. This control algorithm obtains feedback from the UAV localization module and the UAV internal compass in order to actuate any errors in placing the drone due to its inertia.

The AR Drone's motion is controlled by the AT commands protocol. The AT commands are text based commands that are translated by the AR Drone's firmware to control the drone's roll, pitch, gaz, and yaw to specify the direction, speed, and duration of the drone's motion. The controller converts the newly obtained objectives and the current location information to the corresponding set of AT commands. This conversion could be made locally on the drone (i.e. autopilot mode) or on the central node (i.e. remote control mode).

2.5 CPS Optimization

The *UAV CPS Optimizer* module is responsible for updating the objective (i.e. location and orientation) of the drone after obtaining updates on the state of the drone and the environment from the drone's sensor, other drones and the central node. This module has three modes of operation: 1) distributed mode, 2) central mode and 3) emulation mode. In the distributed mode, a distributed objective function is used where each UAV calculates its own new objectives after each update which is suitable for lightweight distributed algorithms. In the central mode, the *UAV CPS Optimizer* reports the updates to the *Central CPS Optimizer* which calculates the new objectives based updates from all UAVs. Finally, in the emulation mode, the UAVs implement a distributed objective function but off-load its calculations to the central node to meet real-time responsiveness requirements.

2.6 Control and Coordination Networks

One of our design choices is to separate the intercommunication between the UAVs and the communication between each UAV and the central node to two different networks. This design choice was motivated by two reasons: 1) While communication between UAVs can be established in ad-hoc,

unreliable fashion, communication between the UAVs and the central nodes should be using an infrastructure that is connected to the UAVs at all times, 2) AR Drone 2.0 [5] provides an SDK for monitoring and controlling the UAVs that requires continuous communication between the UAVs and a central node. We rely on that SDK to facilitate the different tasks of the *Central Node Modules*.

For the implementation of the *Control Network*, we use the built in WiFi card but changed its configuration to work as a client instead of as an access point which is its default configuration. This way all nodes can connect to the central node through a WiFi infrastructure. As for the *Coordination Network*, we plug a WiFi dongle into the drone, compile the dongle's driver to work on the ARM-based Busybox, and install the driver on the drone. The dongle is then configured to work in Ad-Hoc mode and join the drones SSID as soon as its visible.

3. CASE STUDY

In this section, we demonstrate the steps of bringing a sample CPS simulation to life using *UnA*, namely a mobile-camera-based surveillance system. The system's objective is to maximize the number of targets covered by a set of mobile cameras. The proposed algorithm is a central algorithm whose inputs are the location of all targets and the mobile cameras, and its output is a set of location and orientation directives for all mobile cameras. For the rest of the section, we will show the different steps of implementing the proposed CPS using *UnA*.

3.1 UAVs and Targets Localization

To localize the different moving objects in the area of interest (i.e. targets and UAVs), each drone is tagged with a distinct color. On the other hand, because we do not care about the identity of the targets, all targets are tagged with the same color. An Axis 213 PTZ Network Camera is mounted at 5.2 meters to locate the different objects. The central node queries the camera for images at a rate of 20 Hz and uses OpenCV to search each image for all tags by iteratively performing the following steps for each color tag:

1. Filter the image based on the color range of the tag (`inRange()`). The tag's color range is extracted by initial calibration.
2. Filter noisy parts of the image that fall within the same color range as the tag by eroding areas smaller than 3×3 pixels (`erode()`). By this step, the only remaining white area within the image is the tag we are trying to localize.
3. Detect contours using [10] to locate the tag within the image (`findContours()`). Pixel coordinates can then be scaled to actual x,y coordinates.

3.2 CPS Implementation

Our initial simulations results for the coverage algorithm were implemented on Matlab to benchmark the performance of the coverage algorithm using a different set of parameters (e.g. number of targets, number of mobile cameras and targets distribution). The Matlab code was used with a minor modification of obtaining its input from the localization module and sending its output to the UAV controller through sockets.

²While multiple color tags are used in [8] to track the drone's location and orientation, our proposed approach uses a single tag to track the location and the internal campus of the drone to detect the drone's orientation.

³ While PTZ cameras can cost several thousand dollars, MCS costs up to tens of thousands of dollars.

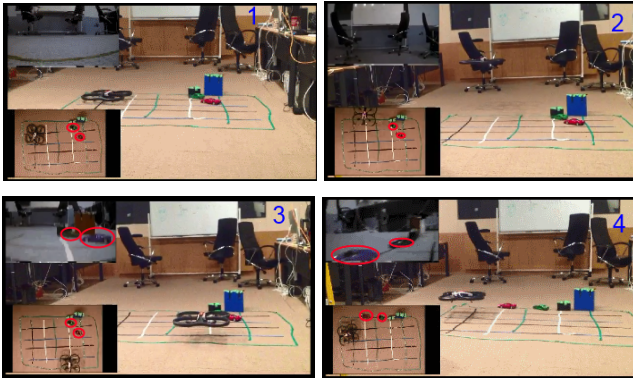


Figure 2: The top left corner is the view from the UAV while the bottom left corner is the view from the master camera. Targets are circled in red when they appear in the view of either cameras.

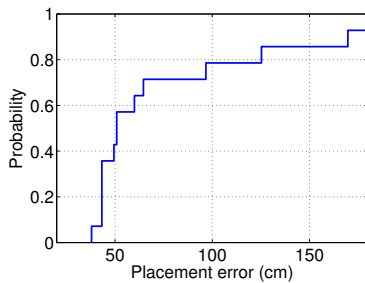


Figure 3: CDF of distance error in drone placement.

3.3 UAV Control

To simplify the case study, all UAV controlling operations are off-loaded to the central node. The central node uses AR Drone API v2.1. to control the motion of the UAVs. The SDK provides the following command `ardrone_tool_set_ui_pad_start` for take off and landing and `ardrone_tool_set_progressive_cmd` for flying and hovering by controlling the UAVs yaw, gaz, roll and pitch. The controller is implemented following the motion methodology presented in Section 2.4.

3.4 Experimental Scenario

The area of deployment was 1.25 by 2.1 meters (4.1 by 6.8 feet). Figure 2 show snapshots of an *UnA*. The scenario of this experiment is to first have the two targets (circled in red) positioned in a certain configuration and have the UAV cover them. Then, move the targets and make the UAV reposition itself to cover the targets.

The first two snapshots show the UAV moving itself to cover the targets in their initial configuration. The third snapshot shows the UAV covering both targets (i.e. both targets are within the view of the drone's camera). The fourth snapshot shows the UAV after moving to cover the drones on their second configuration.

4. LIMITATIONS AND FUTURE WORK

UnA reduces the cost and overhead of CPS experimental deployments, however these advantages does not come with-

out some limitations. One of the major limitations of *UnA* is the accuracy of the UAV placement. Errors in UAV control and maneuvering inherently appear due to inaccuracies in computer vision based localization system and assisted maneuvers provided by the AR Drone's firmware. We benchmarked the *UnA*'s placement error by placing the drone arbitrarily and then moving it to a certain location over 14 times. Figure 3 shows the CDF of the error in placing the drone (i.e. the distance between the drone's actual location and its intended location). The effects of these placement inaccuracies are significantly increased due to the small size of the deployment area. Nonetheless, the current state of *UnA* could be useful for many applications.

Our future plans for *UnA* include larger deployments with two or three drones. We also plan to deploy Click Modular router on the UAVs to allow for a more extensive evaluation of communication protocols between the UAVs. Furthermore, we plan to extend the "physical" components abstraction to allow the seamless integration of several other mechanical systems (e.g. ground vehicles). We also plan to explore other approaches for UAV tracking (e.g. RFID tracking).

5. REFERENCES

- [1] G. Bradski and A. Kaehler. *Learning OpenCV: Computer vision with the OpenCV library*. O'reilly, 2008.
- [2] T. X. Brown, S. Doshi, S. Jadhav, and J. Himmelstein. Test bed for a wireless network on small uavs. In *Proc. AIAA 3rd "Unmanned Unlimited" Technical Conference, Chicago, IL*, pages 20–23, 2004.
- [3] C. Fok, A. Petz, D. Stovall, N. Paine, C. Julien, and S. Vishwanath. Pharos: A testbed for mobile cyber-physical systems. *Univ. of Texas at Austin, Tech. Rep. TR-ARISE-2011-001*, 2011.
- [4] D. Jung, E. Levy, D. Zhou, R. Fink, J. Moshe, A. Earl, and P. Tsotras. Design and development of a low-cost test-bed for undergraduate education in uavs. In *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*, pages 2739–2744. IEEE, 2005.
- [5] T. Krajník, V. Vonásek, D. Fišer, and J. Faigl. Ar-drone as a platform for robotic research and education. In *Research and Education in Robotics-EUROBOT 2011*, pages 172–186. Springer, 2011.
- [6] S. Lupashin, A. Schödl, M. Sherback, and R. D'Andrea. A simple learning strategy for high-speed quadcopter multi-flips. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 1642–1648. IEEE, 2010.
- [7] N. Michael, D. Mellinger, Q. Lindsey, and V. Kumar. The grasp multiple micro-uav testbed. *Robotics & Automation Magazine, IEEE*, 17(3):56–65, 2010.
- [8] H. Oh, D.-Y. Won, S.-S. Huh, D. H. Shim, M.-J. Tahk, and A. Tsourdos. Indoor uav control using multi-camera visual feedback. *Journal of Intelligent & Robotic Systems*, 61(1-4):57–84, 2011.
- [9] R. R. Rajkumar, I. Lee, L. Sha, and J. Stankovic. Cyber-physical systems: the next computing revolution. In *Proceedings of the 47th Design Automation Conference*, pages 731–736. ACM, 2010.
- [10] S. Suzuki et al. Topological structural analysis of digitized binary images by border following. *Computer Vision, Graphics, and Image Processing*, 30(1):32–46, 1985.
- [11] M. Valenti, B. Bethke, D. Dale, A. Frank, J. McGrew, S. Ahrens, J. P. How, and J. Vian. The mit indoor multi-vehicle flight testbed. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 2758–2759. IEEE, 2007.